

ABSTRACT

A method and distributed scheduler for use therewith has at least two clusters of source port modules, each tracking all queues associated with a respective input-node and relating to a respective subset of available input-nodes. Each source port module receives available output-nodes, and generates a weight
5 for each queue therein. Each source port module generates at least one request relating to the highest weight serviceable queue. The respective requests of each source port module are accumulated, and for each cluster of source port modules, the request is chosen for which no two requests relate to the same input-node, and for each output-node, the chosen requests have highest weight. The highest
10 weight request from all clusters is determined in respect of each output node receiving requests from one or more input nodes. A grant is sent to the input-node having the highest weight request.